

# A critique of the philosophical foundations of Effective Altruism

“Making the universe great again!”

~a totalist utilitarian

Bill and Melinda Gates are doing it. Warren Buffett is doing it. Elon Musk is doing it. Trump supporter Peter Thiel does it. The late Derek Parfit (whose “repugnant conclusion” was featured at our last meetup and figures below again) did it... Some, but not all, virtue ethicists ([Justin Oakley](#)) are doing it. Some, but not all, non-consequentialists are doing it ([Larry Temkin](#)). Even—surprise—[some Christians](#) are, too... What would Jesus do? [I wonder....](#)

It all started in 1972 with the publication of Peter Singer’s article [“Famine, Affluence, and Morality.”](#)<sup>1</sup> Singer begins the piece with: “As I write this, in November 1971, people are dying in East Bengal (now Bangladesh) from lack of food, shelter, and medical care.” By 1974, 1.5 million people had died from [a famine](#) that was largely anthropogenic. Too many local and foreign power brokers, the U. S. not least among them, caused the price of food grains to skyrocket, a fateful blow to a subjugated and devastated population.

Singer’s article got anthologized in many introductory philosophy texts, insuring that many thousands of college students whose first exposure to philosophy acquainted them with a simple utilitarian recipe for discharging the moral responsibility each of us has to make this a better world in virtue of being one sentient being among others. His core principle is:

if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it.

The article contains one of the most famous thought experiments in recent philosophy:

...if I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.<sup>2</sup>

If the child were on the other side of the planet and there was something I could do the cost of which, to me, was of relatively little moral significance, then I don’t have an excuse for not doing it. And there is such a thing I might do from wherever I happen to be: I might give a

---

<sup>1</sup> *Philosophy and Public Affairs*, vol. 1, no. 1 (Spring 1972), pp. 229-243.

<sup>2</sup> *Ibid.* See [this short video](#) on Singer’s drowning kid.

certain percentage of my income which would easily go a long way toward helping prevent the suffering and death of children on the other side of the planet. If the child drowning in the pond close by has a moral claim on me, why not the many starving children far away? What moral difference does the *distance* separating us make? (Distance means nothing morally to utilitarianism. Understand this and you understand something critical about utilitarianism.) None that a utilitarian like Singer can see. Sparked by his principle, now fifty years later, an eleemosynary movement has caught fire called Effective Altruism. Who said philosophy and money don't mix?... Here Singer explains the idea in [a TED talk](#). Briefly, EA is guided by the belief that we ought to

1. give much more of our income to charitable causes than we currently do, and
2. that we should marshal every technical means at our disposal to promote the efficient use of the resources we give.

*Give more wisely*, in a nutshell. The end goal is to maximize the total well-being of people and sentient beings everywhere. What's not to like about Effective Altruism?

## Squinting at EA

It is not the implementation problems with EA that will be our primary focus here. There are plenty of those, but there's nothing inherently fatal about them. Sure, people are basically selfish and short-sighted, but most also come equipped with some natural empathy. Singer does no wrong in trying to harness what empathy we have to reason-driven effectiveness. Even if his moral theory turns out to be misguided, you don't have to be a brilliant billionaire or even philosopher to sense that Singer is right to urge his two-point program on a world of moral slackers.

Again, it is the underlying moral theory that underpins the movement that will be our main concern. Are the consequentialist ethics—more specifically, the utilitarian moral model—that drives the movement adequate to the task of making this a better world?—or even the more modest goal of minimizing suffering?

No one that I am aware of questions the intent of EA. (But “intent,” note this, confers no intrinsic moral merit in the eyes of a utilitarian.) No halfway decent person can argue with the goal of helping relieve suffering and doing it with our heads screwed on straight. What is questionable is Effective Altruism's understanding of what is wrong with the world and how to

fix it. Today, organizations, clearly inspired by EA, such as the Bill and Melinda Gates Foundation are acting as filters for efficient philanthropy where many billions of dollars are caught up in a bottleneck, [just waiting for “effective” causes to pass muster...](#)

“Giving away vast amounts of money is way harder than it looks...” says David Callahan, writing at InsidePhilanthropy.com. Here he speculates on [why Melinda and Bill are way behind on getting rid of their money:](#)

### **They Want That Big Fortune to Get Bigger**

This is my favorite theory: Bill Gates wants to keep growing his fortune and have an even bigger pile to pump into philanthropy in a decade or two. Beyond his continued large holdings of Microsoft stock (which might one day recover some of its past valuation), Gates has made a wide variety of investments. It’s perfectly plausible that that \$77 billion fortune could grow even larger—a lot larger. Could we imagine Bill and Melinda Gates sitting on \$150 billion a decade or two from now? Yes, we could.

Regardless, here’s the bottom line: As much as the Gates Foundation is now giving, it’s actually a near certainty that *we ain’t seen nothing yet.*

Really? And in the meantime, there’s all this suffering that might be alleviated. But if in the end more gets alleviated by waiting?...well... maybe so.

Except, again, that there are deeper problems with the idea of Effective Altruism, both at the level of implementation and at the level of theory. For a scathing critique at what happens on the ground at the site of “effective” charity, check out [this podcast](#) from Carneades.org by someone with hands on experience as an aid worker in Africa as well as philosophical skills. Since this is a philosophy forum, my focus here we be on the failings of the theory.

The cause of giving wisely to reduce suffering in the world is overdetermined. There are many reasons why we should give until it hurts. But one of the first things that should hurt is your brain, trying to figure out how to do it best. That has never been easy. The EA movement is to be commended for at least bringing that sometimes forgotten fact to the fore. But its proffered remedy is underwhelming in an especially revealing way.

## Utilitarianism - an extreme sport

Consequentialist theorists like Hilary Greaves<sup>3</sup> have offered interesting arguments for why “total” utilitarianism, as opposed to “average,” “critical level,” or “person affecting” varieties of utilitarianism (about which more below), is the most coherent form of the theory.

But “totalism,” it turns out, makes some pretty amazing demands on us. We are not only to mitigate existing suffering but to expand opportunities for more suffering(!) so that we can intervene with efforts at its mitigation or eradication. What is the answer to the question Parfit was rummaging about so hard looking for when he stumbled on the repugnant conclusion? He was looking for the answer to the question what is the most morally defensible population size? True, population ethics demands we *momentarily* keep our numbers manageable. We don’t want to threaten our very existence. Premature self-annihilation is bad.<sup>4</sup> We have a mission to complete. But, in the final analysis, the answer to Parfit’s query on the morally defensible size of our population is that it *should* be *infinite*. That’s right, as many folks as an infinite universe can contain. We need to harness the universe to support a never-ending augmentation of the numbers of sentient natural kinds. We must populate every corner of the universe. Fill up those friendly rocky exoplanets first, then see what we can do to exploit the resources of the less hospitable ones to *make* planets more friendly to us... in a responsible fashion, of course, we shouldn’t make a mess of it along the way, but it is our ultimate destiny—to *maximize happiness everywhere for all time by all means necessary!*<sup>5</sup>

The morally imperative procedure, if “total” utilitarianism has it right, goes something like this:

- Step 1: minimize suffering for existing people—try to make them as happy as possible,
- Step 2: create more people with an existential risk of suffering, but
- Step 3: nip it in the bud: minimize their risks of suffering by making them as happy as can be by continuing to keep our capabilities a step ahead of our needs.

---

<sup>3</sup> See “Resources” below.

<sup>4</sup> Though Hilary Greaves suggests that, in cosmological terms, the universe may implode at some point, scrunching us up along with it. We won’t argue. With a background in physics, as well as a philosophy, she knows her stuff on this point.

<sup>5</sup> Happiness is like a fluid. We are like containers. Fluid requires a container to prevent its dispersion. An infinite quantity of fluid requires an infinite number of containers—or one very big one. Of infinite size. Or so we must speculate when we try to wrap our minds around the viral idea of “total” utilitarianism.

More is always better if you are a *positive* utilitarian. The best return on our existential investment is to conquer the universe. Before God dropped the ball and let suffering happen, the universe was great. Let's make the universe great again!

Now go back and read [the writeup on David Benatar's negative utilitarianism](#) for relief if all this utility maximizing activity makes you tired. Otherwise, read on to see how utility got to be so arduous, causing even billionaires to take desperate measures...

## Deeper into population axiology

Population ethics urges us to consider large scale and long term implications of our choices to reproduce by imagining the cascade of possible long term effects. Population axiology, as it is also called, is usually couched in terms of consequentialist ethics, specifically utilitarianism. More so than other ethical theories, utilitarianism lends itself to economic styles of fact appropriation. Positive utilitarianism,<sup>6</sup> in particular, concerned with maximizing well-being or happiness, can be applied to population ethics in several ways. The two major ones are Totalism and Averageism. One way to distinguish them is to consider how each would meet the possibility of the extinction of sentience. That would mean the end of beings like us who can experience pleasure, pain, suffering, and happiness.

### *Totalism*

Totalists want to maximize the total amount of happiness ever experienced over all time by sentient species, that is, the sum of all the happiness ever experienced, past, present, and future. Our species, or the set of all sentient species of which we are part of, are assumed to have a definite interest in prolonging life at the highest level of positive happiness, both by maximizing it for the individual case and/or by maximizing it for the size of the set of potential happiness containers. The more containers—such as you or me or other sentients—the more total volume of happiness is possible. More is desirable because we are supposed utilitarians. So... Extinction would be a very bad thing from this perspective. The total amount of happiness is summarily curtailed by extinction.

### *Averageism*

---

<sup>6</sup> For the negative utilitarian view see [our discussion of David Benatar](#). The difference between negative and positive perspectives on the question of the size of a morally justifiable population could not be more extreme: Benatar thinks utility would have been better served had there *never* been sentient beings in existence at all, and given that there are, we should stop procreation in its tracks. Totalist arguments, on the other hand, seem to point to an infinite population of sentients.

According to this view, we should maximize the average well-being level of all the people that have ever lived. This view would only care about extinction if, because of extinction, people whose well-being would have raised the average had they existed, in fact, never come into being. If future people are at a worse level of well-being than the current average, we lose nothing by becoming extinct as a species. Their non-existence in itself is not to be thought of as a loss of something that we should morally concern ourselves with. Remember, these people don't exist now so can't themselves be said to have lost anything. On the other hand, if we can expect them to be better off than us, then there is every reason to lament our extinction: for their non-existence would leave us *as a sentience-equipped species* with a lower average well-being over the full span of our collective existence.

Again, vis-a-vis extinction:

Averagism is only conditionally concerned with extinction: if and only if extinction would lower the average happiness of each being who ever lived.

Totalism is much more concerned with extinction than averageism since any curtailment of the increase of more containers of happiness in the world would be a loss. Any curtailment of lives that are above the threshold below which life would not be worth living lowers the total quantity of happiness that there might otherwise have been.

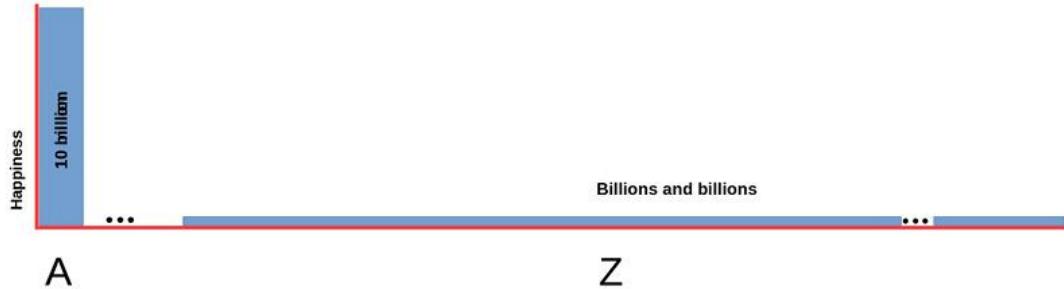
But averagism suffers from the *sadistic* conclusion and totalism from the *repugnant* conclusion.

The repugnant conclusion we are already familiar with from Parfit's thought experiment discussed at [our last meetup](#). But, briefly, it recommends that a world comprised of a vast population of near-miserable-but-not-quite-wretched-beyond-redemption people would be a better world than a world with a much smaller population of the supremely happy. Or, as Parfit put it, we are led to:

*The Repugnant Conclusion:* For any perfectly equal population with very high positive welfare, there is a population with very low positive welfare which is better.<sup>7</sup>

---

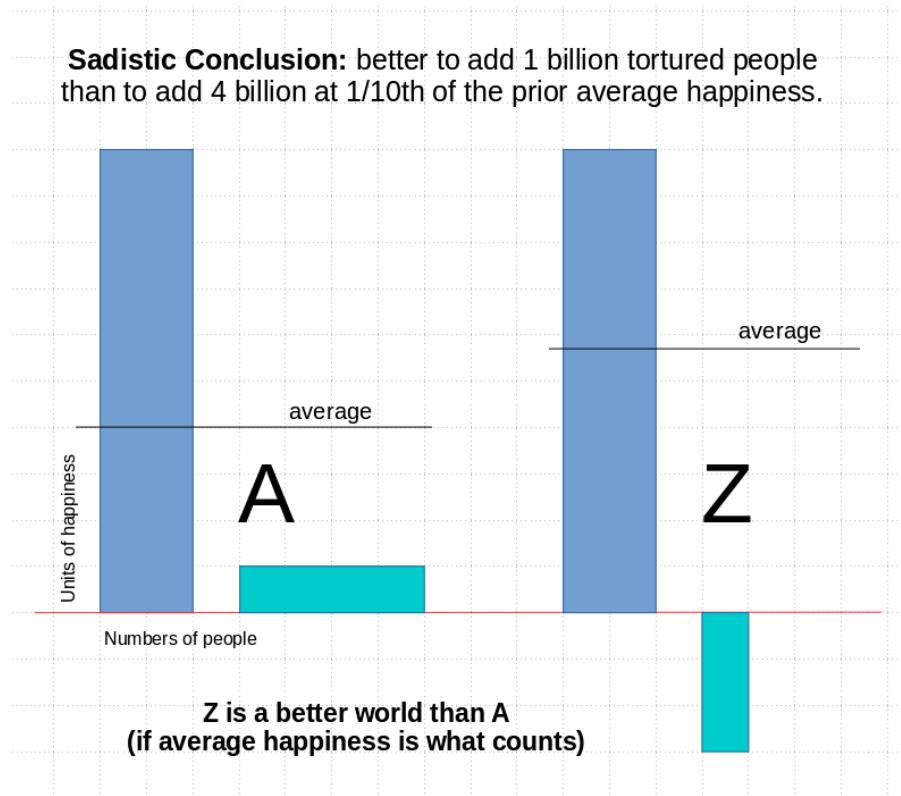
<sup>7</sup> Derek Parfit, *Reasons and Persons*, 1984, p. 388.



Z world, the repugnant conclusions contends, is better than A world. Measure the area in the blue rectangles to see why. Height is happiness level. Width is population size. If *total* happiness is your goal, you have to face this outcome.

However, you can escape this conclusion if you shift your priorities and place *average* happiness, not total happiness, as your target for maximization. Z world clearly does not have higher average happiness than A.

But the problem with this move is the Sadistic Conclusion:



By focusing on average happiness, we leave on the table the possibility that we may generate a world with a higher average happiness level relative to an acceptable world by adding a

smaller number of people with a much lower happiness quotient. In other words, a world with a small population in a hellish state (Z) might be preferable to a world with more people in a much better state (A).<sup>8</sup>

Other ways of handling population axiology have been proposed but they seem subject to even worse problems. The Critical Level theory, for instance, attempts to set a level of well being above which adding more happy people is fine and below which it is not. The trouble is that such a solution is subject to *both* the repugnant and the sadistic conclusion. In essence, all it does is raise the bar of what would be acceptable while leaving room for the machinations that generated those impasses to be set in motion again: only at a higher level of what is better and worse.

Another theory, Person-Affecting theory, attempts to limit our concern to only the happiness of existents and ignore that of merely hypothetical people. After all, some hypothetical people are of negligible concern to us: no doubt, those whose existence would not begin until many generations from now. But others such as descendants we may know or can, with a little imagination, conjure—grand-descendants, and great-grand-descendants—though hypothetical, are of concern to many. It's not clear we are willing to say their well-being can be dismissed out of hand.

Hilary Greaves sums up the situation:

In his seminal discussion of population ethics, Parfit (*[Reasons and Persons]*1984) rejected Totalism and Averagism for reasons including those given above, failed to find any alternative axiology that he himself considered satisfactory, but held out hope that this was merely for want of searching hard enough: that, in the future, some fully satisfactory population axiology,<sup>9</sup> called ‘Theory X’ by way of placeholder, might be found. Much of the subsequent literature has consisted of attempts to formulate such a ‘Theory X’. However, as we have seen above, every extant population axiology is open to serious objection: if it does not entail the Repugnant Conclusion then it entails the Sadistic Conclusion, or is anti-egalitarian, or has obviously unacceptable implications concerning future-people cases, or otherwise leads to some similarly serious objection.

---

<sup>8</sup> The picture here reminds us of why some historians have argued that slave economies were, in fact—especially if you drop moral scruples and add some technological sophistication (not things that necessarily go together)—rather effective. See “[Did slavery make economic sense?](#)”, C.W. and A.J.K.D, *The Economist*, Sep 27th 2013.

<sup>9</sup> The “ordering of states of affairs in terms of better and worse overall,” as Greaves defines it.

Greaves goes on to discuss the view developed by philosophers in population axiology that there cannot be a “Theory X”—ever—because of certain “impossibility theorems” which seem to suggest that the underlying intuitions that everyone wants to hold *cannot* be held simultaneously. We want impossible things. Something has to go.

Still another way out, on its face no more savory, is to just *accept the repugnant conclusion*. The reason it makes us uncomfortable, right off the bat, is that we haven’t been properly acculturated to the idea. Give us a little time and eventually our hesitance with the notion that vast numbers of equally, but *just barely*, happy people is indeed where ethical thinking should take us will fade. We’ll adapt. Right?

## Can we fix totalism?

Can we do better? Driven by the evolutionarily sanctioned optimism that no doubt got us here,<sup>10</sup> we may take another desperate look at the *totalist* version of utilitarianism (the others seemingly too bleak to pursue). We consider what might be done to steer it clear of the repugnant conclusion while still offering direction in the project of finding an answer to the question what population number is defensible on utilitarian moral grounds.

The obvious way is to add people<sup>11</sup> who are *never* at a happiness level below the current one. To be clear, this *rules out*:

1. extinction (except through cosmological disaster not in principle avoidable by us, e.g., the implosion of the universe), and
2. holding the population steady indefinitely at any number (because this violates totalist utilitarianism).

Since we are not allowed to stop or go backwards, onward! That leaves us, as our moral teleology, the goal of *expanding sentience without end*. We have a moral imperative to add every last bit of sentience the universe will support.<sup>12</sup> This entails a morally driven, technological mandate to pursue any and every means to these ends:

---

<sup>10</sup> But [check out Benatar](#) on the likely possibility that rose-colored glasses (aka, the “Pollyanna Principle”), however critical to survival, may not be epistemically warranted, however morally driven.

<sup>11</sup> For present purposes, we use “people” interchangeably with the more inclusive “sentient beings.” This is a standard utilitarian presupposition.

<sup>12</sup> Totalism is perfectly compatible with strategic population control. If, in the long run, we will be fruit indefinitely, and to do that we need to hold back, even step back, our numbers *momentarily*, that is perfectly in order. We may need to pace ourselves on the way to conquering the universe.

1. adding more and more people, and
2. more and more non-human forms of sentience, and
3. making sure every generation of sentient beings is hedonically better off than the previous one (pauses and reversals justifiable *only* as conducive to a general progressive trend).

These ends might be served by directing our best efforts to:

4. Ending the suffering of the present population (enough of that already!).
5. Seeking to colonize other planets (get on with colonizing the moon, Mars, and beyond).
6. Seeking to survive for all time (make aging optional and mandate its avoidance).
7. Anticipating every setback for future populations (don't do the dinosaur thing, have an asteroid plan!).
8. Developing the means to triumph over those setbacks. (R&D—lots!).
9. Enhancing the quality of natural sentience. (Make ourselves ever more sensitive to pleasure and pain? (Or more of one and less of the other?) The possibility of doing this opens up a huge philosophical can of worms! Consider David Pearce's [The Hedonistic Imperative](#) and [Transhumanism](#).)
10. Creating and improving artificial sentience. (Replace ourselves with supersentient robots? Remember: it is the stuff *in* the containers that counts, not the packaging. In conversation with David Pearce, [Singer says](#), “It’s not important to me that it be biological human beings who experience the happiness...”)

Nothing less will do if you are a totalist utilitarian. We fail morally to do anything else. It is not surprising then that *existential risk*—the idea that sentience might come to end, that all beings capable of happiness might someday vanish—is a major preoccupation of some totalist utilitarians. The fear is not merely that the end might be full of pain and suffering but that the event of extinction will forestall all the future happiness: all of it there *might have been* may not be.<sup>13</sup>

---

<sup>13</sup> Nick Bostrom, for example, [has argued](#) that we have reason to fear that we are not what we think we are. One of three possibilities regarding our current ontological status is that we are computer simulations reenacting the past history of distant descendants so advanced in the creation of virtual realities that we are (1) “living” in a computer game of sorts of their making because they (for some reason) have an interest, however idle, in running the simulation. (Boredom?) Or (2) they haven’t the least curiosity about us: we indeed *are* what we think we are but, in the cosmological scheme of things, just aren’t that interesting. And one of these two possibilities is inevitable, he says, otherwise (3) we do not survive long into the future. We don’t really exist, or we exist but are cosmologically boring, or—our days are numbered.

## Hands dirty and heads clear

What does this mean for the philosophy behind Effective Altruism?

It means EA needs, just for starters, to get its moral act together. It cannot, for example, tolerate:

1. Poor, unrealistic assessments of what actually happens on the ground—at *the place where money meets need* in the interest of generating positive numbers. It cannot think that bug nets are a *sustainable* way of bettering both the quality and quantity of human sentience in sub-Saharan Africa.<sup>14</sup> Effective altruists must get their hands dirty with the complex nature of *actual*—not merely theoretical—human desire and need. They should attend to cogent criticism such as [this from those on the front lines of helping people](#).
2. Philosophical confusion about whether EA wants to increase human life years or reduce them. Effective altruists must get their heads clear whether, for example, they are going to save children’s lives from fatal diseases<sup>15</sup> or whether they are going to support efforts whose most salient effect is the reduction of the numbers of existing children such as planned parenthood. As Hilary Greaves has noticed,<sup>16</sup> *You cannot coherently have it both ways*. Not if you profess—as the intellectual leaders of the movement do—totalist



<sup>14</sup> As Danaher summing up Gabriel puts it, “The poorest of the poor suffer from complex, multidimensional forms of deprivation that are difficult to overcome via any one policy intervention.” Shall we just give up then on the poorest of the poor? As in Gabriel’s TX, forget about illiterate widows and the disabled in rural areas vis-a-vis literate urban men? There’s more bang for your buck helping the latter.

<sup>15</sup> In opposition, say, to [the Garret Hardin solution](#).

<sup>16</sup> Hilary Greaves makes this point in a presentation on counterproductive charities supported by EA: *you can fund family planning or childhood mortality causes but can you do both?* [“Repugnant Interventions - Doublethink in Global Prioritization”](#) (video).

conceptions of utilitarianism.<sup>17</sup> You can't salvage with one hand what you preclude with the other.

And if we drop the totalist utilitarian underpinning of Effective Altruism, what is left that is *remarkable* about the idea of EA?

William Adolphe Bouguereau (1825-1905) - [Charity](#) (1878)

At the practical level, the core ideas of EA are not new. People wanting to relieve suffering in the world have always wanted to do so effectively. Perhaps Effective Altruists are saying, “We have the know-how today to do this better. Well-meaning people in the past were hobbled by ignorance and human vagary.” One can hope this is the case but—except by EA’s own measures—and how objective can those be, confined as they are to the measurable?—even some of the tiny contingent of charities it approves of are not doing well if the folks on the ground at the site of the nominally “effective” charitable acts are to be believed.<sup>18</sup> It seems, at times, difficult to distinguish EA from the slickest marketing campaign ever for bug nets. One need not deny the utility of bug nets to notice something funky here. Moreover, because the elitist standards of its flagship charities do not tolerate much in the way of so-so effort, much of its money is sitting idly, waiting for the rest of the world to catch up with its eleemosynary austerity. This is a problem given Effective Altruism’s stated ends. You cannot even begin to be “effective” if you are not even adequate.

But most importantly, from a philosophical perspective, Effective Altruism is asking more of utilitarianism than it can supply. And without utilitarianism, what can we say about Effective Altruists?

---

<sup>17</sup> Of course, if you are not inclined to utilitarianism, you don’t have this problem. You are likely to have others, but we leave discussion of those for another occasion.

<sup>18</sup> For example, how many mosquito nets are not being used as intended but for more dire needs? How many are implicated in environmental disasters in the making? Apparently, we still don’t know. Where was all the technical know-how behind the Gates Foundation effort *before* this happened? See “[A Tragic Choice: Fight Malaria or Starve](#)” or read “[Meant to Keep Malaria Out, Mosquito Nets Are Used to Haul Fish In](#)” by Ben C. Solomon in *NYT* Jan. 24, 2015. The article reports, “Western governments and foundations donate the money. Big companies like [BASE](#), [Bayer](#) and [Sumitomo Chemical](#) design the nets. They are manufactured at about \$3 apiece, many in China and Vietnam, shipped in steel containers to Africa, trucked to villages by aid agencies, and handed out by local ministries of health, usually gratis.” ([Sumitomo Chemical](#), at least, makes an effort to manufacture in Africa.) Speaking of big companies making billions of dollars abroad in the manufacture and delivery of nets... think for a minute: you give your money to highly vetted “effective” charity, which proceed to fund the nets. Global corporations get the money, Africans get the nets... One would hope that they would get much more, like proper training on how to subsist without cannibalizing nets, which, because they are too fine kill fish by the millions too small to eat, and which contaminate the water with the insecticides impregnated in the high-tech nets.... Once you get the facts straight, giving effectively is supposed to be simple. Just trust experts to get it right. We expect this kind of messiness of old fashioned charities, the kind that give first and think later, not from EA.

“Some, I assume, are good people,” Donald Trump would say.

## Note on theory

I avoid addressing a class of criticisms that (perhaps) begs the question about utility. A great deal of human endeavor does *not* qualify as moral, except instrumentally, by the strictest utilitarian standards. Not art, not the humanities, not mathematics or science, not even time spent on *other* philosophies can claim exemption from the calculating eye of the sentience-minded maximizer. You can’t even just be a money-grubbing Wall Street banker without being shamed into giving a morally significant share of your take to effectively relieve suffering somewhere in space and time, MacAskill suggests. So it isn’t surprising to discover charity should not be any different...

I don’t suggest that other moral theories have it easier deciding the final aims of sentient and rational beings. Deontology, for one, can make demands on “moral” intuitions easily as hyperbolic. But I do claim that utilitarianism *especially* wades through its own set of demands rather artlessly.<sup>19</sup> Again, Singer and company, of course, are right to squeeze what they can from obese camels, heaven or no heaven. But all the theory in the world is no match for human meanness. There isn’t going to be a “Theory X,” *pace* Parfit. Not because it is necessarily impossible, but because *we* are.

## Resources

Peter Singer’s iconic [drowning kid summed up](#) (video). Singer’s seminal 1972 article, [“Famine, Affluence, and Morality”](#) and his 2013 [TED talk on Effective Altruism](#).

[“Charity & Ethics”](#) by Mike LaBossiere (from whom I borrowed the notion of the Bouguereau image).

---

<sup>19</sup> Kant’s moral psychology is much more nuanced and informative than its counterpart in utilitarianism. The highly reductive strategy of the consequentialist program is simply incapable of doing moral philosophy for humans *as we find them*. If the goal, rather, is the *reformation* of human psychological architecture from Baroque to Bauhaus, simplistic, decisive answers to moral challenges begin to emerge that may fascinate our utilitarian tendencies. In which case, the best solution to the question about the size of the morally defensible population of human beings, described as *essentially* sentient, may plausibly be the one suggested by totalism. The population *should* be maximized to infinity—or to extinction, whichever comes first... Whether this is a *reductio* or a *modus ponens*, whether it disqualifies utilitarianism from proper theoryhood or recommends it to the skies, I leave for another occasion.

Iason Gabriel attempts a balanced critique, [“Effective Altruism and Its Critics”](#).

[John Danaher expands on Gabriel](#).

Jeff McMahan defends EA, [“Philosophical Critiques of Effective Altruism”](#).

Rhys Southan, [“Is it OK to make art? If you express your creativity while other people go hungry, you’re probably not making the world a better place”](#). *Real utilitarians don’t bluff!*

Brad Leiter is [somewhat critical of Jeff McMahon defense of EA](#).

Sam Harris interviews William MacAskill, the new youthful torch bearer of EA: [“Being Good and Doing Good: A Conversation with William MacAskill”](#) (podcast).

Richard Yetter Chappell: [“Effective Altruism, Radical Politics and Radical Philanthropy”](#).

Hilary Greaves on the specter of counterproductive charities: *you can fund family planning or preventing childhood mortality but can you coherently—on the assumptions of EA—do both?* [“Repugnant Interventions - Doublethink in Global Prioritization”](#) (video). Greaves [presents at conference on EA](#) (videos). Greaves here [“Extinction risks and Population Ethics”](#) discusses totalism, averageism and person-affected views of utilitarianism and their implication for morally defensible populations (video).

[David Pearce and Peter Singer](#) (video). Suffering? Enough already. David Pearce’s [The Hedonistic Imperative](#).

## Effective Altruism - specific links

[Against Malaria Foundation](#)

[Give Well’s top charities](#)

[Giving What We Can’s responses to criticism](#)

[The Life You Can Save](#) - Peter Singer’s own more expansive EA filter organization.

[GiveDirectly](#) - In “[This Kenyan village is a laboratory for the biggest basic income experiment ever](#)” (VOX.com), Dylan Matthews reports on an EA effort which takes risks with the presumption that we “the givers” are in a position to evaluate “effectiveness.” This charity stands out in its trust of recipients. A such, it’s philosophical underpinning is less consequentialist—or, perhaps, it is *instrumentally* consequentialist, thus escaping many of the criticisms offered here.



The briefest account I know on the proper *attitude* to doing good (or its attempt) is [this passage from Georges Bernanos's 1936 novel The Diary of Country Priest](#). Mind you, I'm not a believer. I recommend it because it hints more accurately at the psychology of good intentions and what it must contend with than anything I've ever read, beside which, the problems of utility are child's play.